ConneXions

The Interoperability Report

December 1989 Volume 3, No. 12

ConneXions — The Interoperability Report tracks current and emerging standards and technologies within the computer and communications industry.

In this issue:

ConneXions is published monthly by Advanced Computing Environments, 480 San Antonio Road, Suite 100, Mountain View, California 94040, USA. Phone: 415-941-3399. Fax: 415-949-1779.

ISSN 0894-5926

From the Editor

The protocol standardization process in the Internet community has been described in previous issues of *ConneXions*. Draft specifications—which may the output from a task force or working group or an individual contribution—are submitted to the RFC editor. After some time, the document may be adopted as an Official Internet Standard. In the OSI world things are very different. Protocols are developed as a result of close international cooperation within the standardization bodies. In general the process takes much longer than in the TCP/IP world, but in fairness OSI is a very ambitious project so one would expect a longer development cycle. Bringing OSI to the end user takes more than the work of standardization bodies. This month, David Chappell describes the various OSI "players" in our series *Components of OSI*.

How many classes of the OSI Transport protocol are there really? Lyman Chapin explains in his Letter to the Editor on page 11.

James VanBokkelen is back with an article on how different kinds of protocols can coexist on a single local area network. The article is reprinted with permission from the FTP Software newsletter called *news@ftp*.

In our July 1988 issue we listed a number of sources for information on TCP/IP. This month we include a much expanded list. The list was compiled with the assistance of the *USER-DOC Working Group* of the Internet Engineering Task Force (IETF). This group is compiling a bibliography comprising documents (both online and hardcopy), reference materials, and training tools addressing general networking information and "how to use the Internet." The purpose of this effort is to provide assistance to end-users and those who provide direct services to end-users. In addition to text documents, the bibliography will also include reference media such as audiotapes, videotapes, and fiche, networking workshops and conferences, and mailing lists. The USER-DOC working group plans to publish both an on-line and hard copy of this bibliography. For more information or suggestions for additions to the bibliography, please contact Tracy LaQuey, tracy@emx.utexas.edu, or Karen Roubicek, roubicek@nnsc.nsf.net.

We hope you have enjoyed *ConneXions* in 1989. As always, your comments and suggestions are very welcome. *ConneXions* will continue to cover topics in the interoperability field in 1990. Have a peaceful holiday!

Components of OSI: A taxonomy of the players

by David Chappell

OSI is the creation of many different people, grouped into many different organizations. The goal of this article is to describe how these organizations work together to turn raw ideas into implementable standards, which can then serve as the basis of interworkable products.

Standard bodies

The most fundamental of all of OSI's documents are the base standards themselves. These standards are produced by a variety of national and international bodies. Three of these—ANSI, ISO, and CCITT—are described here.

ANSI

Most industrialized countries have a national standards organization, responsible for coordinating the development of standards for that country. In the United States, this national body is the *American National Standards Institute* (ANSI). Like most national standards bodies, ANSI produces standards in many areas, ranging from computer languages to screw thread sizes. These standards are produced by *Accredited Standards Committees* (ASCs), each responsible for a particular area.

The ASC with responsibility for computer-related standards is called X3. X3 is in turn divided into several *Technical Committees*, with each assigned a particular area. Among the more important OSI-related Technical Committees are X3S3, responsible for the OSI lower layers, X3T5, responsible for the OSI architecture and upper layers, and X3V1, responsible for office-related standards such as electronic mail. These committees are further divided into *Task Groups* which specialize even further. Most of the U.S. work on the OSI Transport and Network layers, for example, is performed by X3S3.3, while the bulk of U.S. development for the upper three layers of OSI is done by X3T5.5.

Membership in ANSI is voluntary, and is typically drawn from the manufacturing and user communities. As is often pointed out by critics of standards, there are no entrance exams for participants, and no special expertise or experience is required to become a voting member of a committee. And like most standards organizations, positions are arrived at by voting, although the goal is to reach consensus—simple majority is generally not sufficient.

Although ANSI sometimes produces standards on its own, in the OSI world, ANSI committees develop the U.S. position for participation in ISO, described next.

ISO

The U.S. has ANSI, the United Kingdom has the *British Standards Institution* (BSI), and most other industrialized nations have their own national standards bodies. National standards alone, however, are often not enough. If countries can agree among themselves, international standards may be created, adhered to by many nations. The *International Organization for Standardization* (ISO) provides the forum for reaching this multinational agreement. (In spite of the apparently obvious correspondence, "ISO" does not stand for "International Standards Organization.") The members of ISO are the various national standards bodies, such as ANSI and BSI. Like the national bodies themselves, ISO primarily represents the manufacturing and user communities.

Technical committees

ISO is broken into several different *Technical Committees* (TCs). Until a couple of years ago, TC 97 had overall responsibility within ISO for information processing systems. More recently, this function has been combined with a similar function performed by the International Electrotechnical Commission (IEC) to create ISO/IEC Joint Technical Committee (JTC) 1. ISO/IEC JTC 1 is currently the parent TC for OSI work.

Technical committees, in turn, are divided into *Subcommittees* (SCs). Among JTC 1's many subcommittees are several with responsibilities for developing OSI standards. Chief among these are SC 6, responsible for the OSI lower layers, SC 18, responsible for office-related standards, and SC 21, responsible for the OSI upper layers and the overall OSI architecture. (The correlation with ANSI's X3S3, X3V1, and X3T5 committees is no accident. When ISO changes its organization, as it does every few years, ANSI typically modifies its committee structure to align with the new order.)

ISO operates as a consensus-oriented democracy. Each member country has one vote, but the goal is to achieve complete agreement among all interested members. Unsurprisingly, reaching this goal often requires time-consuming discussions and painful compromises.

Document cycle

All OSI standards developed by ISO begin as work items. New work items must first be approved by JTC 1. Next, the appropriate subcommittee solicits contributions from member bodies in the form of *Working Drafts*. These drafts are reviewed and modified, and eventually advance to the status of *Draft Proposal* (DP). The existence of a Draft Proposal indicates at least tentative agreement on the fundamental technical content of a standard.

DPs must be approved by all interested member bodies. To reach this approval, a ballot period (generally 90 days) occurs, with each member body asked to vote. In a typical scenario, some members will return a vote of "no with comments," with the agreement that their vote becomes a "yes" if all comments are resolved. If major technical changes are required, the document may undergo a second round of voting, this time as a second DP. Once all negative votes are resolved, however, the DP (or second DP) becomes a *Draft International Standard* (DIS).

Theoretically, a DIS undergoes only editorial, not technical change. In practice, however, technical changes to OSI documents at the DIS level are not unusual. In some cases, a second DIS may even be required. Eventually, after another round of voting by member bodies, a DIS advances to an *International Standard*. The length of time from new work item to International Standard is variable, but is generally several years.

International Standards describe a final agreement among member countries, and are quite stable documents. Even these final agreements are subject to review, however, and may be modified by the addition of addenda. These addenda pass through three stages: *Proposed Draft Addendum* (PDAD), analogous to Draft Proposal; *Draft Addendum* (DAD), analogous to Draft International Standard; and *Addendum*, analogous to International Standard.

continued on next page

A taxonomy of the players *(continued)*

OSI standards produced by ISO are assigned both names and unique numbers. A number is assigned to a document when it becomes a Draft Proposal and does not change. For example, Draft Proposal 7498 became Draft International Standard 7498, then finally International Standard 7498. Addenda are numbered sequentially, e.g., International Standard 7498, Addendum 1.

In many countries, all telecommunications as well as all mail delivery are run by a government monopoly, usually known as the *PTT* (for Post, Telegraph, and Telephone). A major exception to this is the United States, with its regulated but competitive market. (To get some idea of what a U.S. PTT might be like, imagine the U.S. Postal Service running both Western Union and the pre-divestiture AT&T.) These government agencies wield considerable power within their country.

CCITT

To achieve effective international communications, however, international agreements among the telecommunications providers are required. The forum for reaching these agreements is the *International Telegraph and Telephone Consultative Committee* (CCITT). (The acronym comes from the group's French name, Comité Consultatif Internationale de Télégraphique et Téléphonique.) CCITT is part of the *International Telecommunications Union* (ITU), which is itself a part of the United Nations.

CCITT has several *classes* of membership. The highest class is A members, which includes the PTTs (the United States A member is the U.S. State Department). Next are B members, which include *Recognized Private Operating Agencies* (RPOAs). Examples of RPOAs are private communications companies such as AT&T and MCI. Other classes of membership allow participation by interested organizations, such as manufacturers of telecommunications equipment and other industrial or scientific organizations. Like ISO, CCITT operates as a consensus-seeking democracy, with only A members allowed to vote.

Also like ISO, the actual work of creating these documents is performed by several subgroups, called study groups. These study groups are then further subdivided into working parties. The major CCITT groups involved in OSI are Study Group (SG) VII, with responsibilities throughout the seven layers of OSI, and SG VIII, responsible for certain upper layer issues.

Recommendations

The documents produced by CCITT are called *recommendations* rather than standards, although they serve the same function as do standards. Like ISO standards, CCITT recommendations pass through stages on their way to full acceptance, although CCITT operates somewhat differently than ISO. Every four years, a new set of recommendations is produced, covering a broad range of communications topics. These recommendations are approved by the Plenary Assembly, which also approves the general direction of the work to be done in the next four year study period. The recommendations for a given year are published in a series of books, all with the same color cover. The covers were orange in 1976, yellow in 1980, red in 1984, and blue in 1988. The recommendations for a given year are commonly referred to by color, so, for example, the 1988 CCITT Recommendations are called the *Blue Books*.

During each four year study period, the actual work is accomplished by the Study Groups and Working Parties assigned responsibility for a particular area. These various groups produce Draft Recommendations, which typically become Recommendations after the Plenary Assembly.

Four years can be a long time to wait for a new recommendation. When agreements are needed more quickly, accelerated procedures may be invoked, allowing interim approval of recommendations between Plenary Assembly meetings

CCITT Recommendations are assigned names and alphanumeric designators. These designators consist of a single letter, a period, and a number, such as V.24 or X.25. These designators are usually retained throughout the life of a recommendation; CCITT Recommendation X.25, for example, appears in slightly different versions in the orange books, yellow books, red books, and blue books.

Cooperation between ISO and CCITT

ISO and CCITT have, for the most part, worked together throughout the development of OSI. (In fact, some important parts of OSI, such as X.400, were originally created almost entirely by CCITT. For this and other reasons, referring to OSI as "the ISO protocols" is a misnomer.) This voluntary cooperation is itself a notable thing, coming as it does between two standards-making organizations of approximately equal stature. Also, the people active in ISO and CCITT tend to come from quite different perspectives: ISO members from the computer industry, CCITT members from the PTTs and RPOAs. The differing goals and perspectives of these two groups has sometimes led (and will no doubt continue to lead) to heated discussion and difficult compromise.

It may also lead to some confusion about the documents produced by the two groups. ISO and CCITT commonly publish nearly identical text as both an International Standard and a Recommendation. For example, ISO publishes the protocol for the OSI Transport layer as International Standard 8073, while CCITT produces technically equivalent text as Recommendation X.224.

The people who actually do the work in both ISO and CCITT, and in the national bodies on which they depend, are volunteers drawn from government, academia, and industry in many countries. The OSI standards they have produced, hundreds of documents which often describe very complex protocols, are the result of literally millions of hours of work. Their creators must solve not only difficult technical problems, but also the often more difficult organizational problems attendant in creating worldwide standards. Seen in this light, the achievements of OSI are as much political as they are technical.

Implementors Workshops

The creation of OSI standards is important. Without the standards, there would be no protocols; without the protocols, there could be no actual implementations and thus no products. But standards alone are not enough. By their very nature, they typically suffer from several deficiencies that may prevent building actual implementations which can communicate with one another.

First, standards are written in human languages, e.g., English. Human language is inherently imprecise, and the text of an OSI standard is no exception to this rule.

continued on next page

A taxonomy of the players *(continued)*

For two implementations of the same protocol to correctly communicate, both implementors must interpret the standard in exactly the same way. With OSI, where those implementations will be done by different people, on different computer systems, possibly in different countries, a single interpretation of the standard becomes essential.

Second, human beings, even the creators of OSI, are imperfect. They occasionally make mistakes and forget things. Many (perhaps most) of the OSI standards were finalized well before any implementations based on the protocols defined in those standards existed. As a result, errors and omissions exist in the standards documents.

Finally, the standards which make up OSI typically contain *choices*. The consensus orientation of the standards bodies makes this almost a certainty; what better way to resolve an argument than by accepting both positions? Unfortunately, while this allows the creation of a single standard, it does not necessarily allow the construction of interworkable implementations. Two implementors can both follow the same standard and yet, if each chooses different options within that standard, they may still be unable to communicate. For OSI to work, these problems must be resolved.

NIST OIW

All of these problems appear when implementing OSI. Therefore, it makes sense that groups of implementors should form to deal with these issues. In the U.S., the *National Institute of Standards and Technology* (NIST, formerly called the National Bureau of Standards, or NBS) sponsors an *OSI Implementors Workshop* (OIW). At these quarterly meetings, people involved in the actual implementation of OSI protocols (generally not the same people who created those protocols) agree on interpretation, problem fixes, and protocol choices. (With very few exceptions, the "base" ISO and CCITT standards are not modified by the workshop.) The agreements reached at these workshops can then form the basis for creating interworkable implementations.

The NIST OIW is divided into several *Special Interest Groups* (SIGs), each responsible for implementation agreements on some particular topic. SIGs currently exist for several areas, including X.400, FTAM, VT, Network Management, X.500, Lower Layers, and Upper Layers. Each SIG develops agreements for its area, which must then be approved by the Workshop plenary, consisting of all Workshop attendees. Initial agreements become part of the Working Implementation Agreements, while once a year (so far, each December), all completed agreements are published by NIST as *Stable Implementation Agreements*.

EWOS

In Europe, similar tasks are performed by the *European Workshop on Open Systems* (EWOS), while the *Asian and Oceanic Workshop* (AOW) carries out the same function for Japan, Australia, and the Pacific Rim countries. This multiplicity of implementors groups raises a significant problem: what happens if the groups reach different agreements? In an attempt to solve this problem, ISO has created a new class of documents, called *International Standardized Profiles* (ISPs). The intent is first to allow the three regional workshops to harmonize their implementation agreements among themselves, then to have ISO publish the result as an ISP after a short ballot period.

This ISP will provide a single worldwide implementation agreement, allowing implementors in all countries to build interworkable implementations. (Whether this process will actually work remains to be seen; the first ISP, for FTAM, is currently being balloted.)

A fairly obvious question arises: why bother with this two-tiered system of standards followed by implementors agreements? Why not just have the standards bodies do the whole thing? Maybe the most correct answer is "because the current system seems to work." OSI standards bodies tend to be populated with somewhat abstract, general thinkers, especially in the upper layer committees. While this is probably appropriate, since the goal is to define standards useful for a broad range of purposes, many of the main contributors to OSI standards do nothing but develop standards, i.e., they are not involved in actual implementations of the protocols they design. The implementors workshops, on the other hand, tend to be made up of people who are actually building OSI products. Their perspective is somewhat different, tending more towards "How can we make this standard implementable so I can meet my product delivery deadlines?" Applying this second approach to the results of the first, i.e., to the ISO and CCITT standards, seems to yield acceptable results.

Profile groups

ISO and CCITT produce the base standards for OSI. Implementors workshops further refine those standards into a better base for real implementations. But even the implementors agreements contain some choices. For example, the NIST Stable Agreements contain specifications for more than one class of Transport protocol as well as agreements for a large number of OSI application service elements. So while these agreements do narrow the choices somewhat, building truly interworkable products requires one more thing: conformance to a particular profile. A profile specifies exactly which protocols are required for a particular environment. Two currently popular profiles (also called stack specifications) are those defined for MAP/TOP and for GOSIP.

MAP/TOP

The first widely-known OSI profile was the *Manufacturing Automation Protocol* (MAP), first promulgated by General Motors. Consisting of OSI protocols appropriate for the factory environment, the banner of MAP has since been taken up by a larger group of companies. For the most part, MAP is based on the NIST OSI Implementors Workshop Stable Agreements. In some cases where OSI standards have not reached completion, such as network management, MAP includes protocols based on preliminary standards work.

MAP is intended for the factory floor. For office-based OSI communications, Boeing Computer Services developed the profile called *Technical and Office Protocol* (TOP). TOP is quite similar to MAP, with differences primarily in the bottom two layers and in the Application layer. Like MAP, the development of TOP has spread from its original sponsor to include a significantly larger group of companies.

Exactly which protocols are included in MAP and TOP are determined by one of the many MAP/TOP committees. The MAP and TOP profiles have both had more than one version; the current version of each is numbered 3.0.

continued on next page

A taxonomy of the players *(continued)*

GOSIP

The *Government Open Systems Interconnection Profile* (GOSIP) specifies exactly what must be implemented to sell OSI products to the U.S. Government. Like the MAP/TOP specifications, GOSIP is based on the NIST OSI Implementors Workshop Stable Agreements. Unsurprisingly, then, GOSIP has much in common with these earlier profiles. It is produced by a committee of representatives from various divisions of the federal government, and is published by NIST. The first version of GOSIP was published as *Federal Information Processing Standard* (FIPS) 146 in 1988, and new versions are planned for coming years. Compliance with GOSIP will become mandatory for many federal procurements beginning in August 1990.

Other countries, such as the U.K., have their own versions of GOSIP. To distinguish among these national versions, the country name is often used, e.g., U.S. GOSIP or U.K. GOSIP.

COS

Without standards, no OSI protocols would exist. Without implementors agreements and profiles, those standards would remain too imprecise to actually implement. But there are still other things needed before the universal interconnection promised by OSI can become a reality. The *Corporation for Open Systems* (COS) was created to address some of these remaining issues

First announced in December 1985, COS is an international consortium of companies, each of which contributes some amount of money towards the support of the organization. Among its members are IBM, AT&T, and DEC, as well as many smaller companies. The primary goal of COS is to promote OSI and ISDN. COS neither creates standards nor produces implementors agreements. Instead, it adopts already created standards and agreements and assists in making them a reality. Among the organization's major contributions has been the development of another piece of the OSI pie: *conformance tests.*

For implementations built by different vendors to interwork, all must strictly obey the rules for behavior defined by each of the OSI protocols. But how can the purchaser of an OSI product know that the product in fact conforms to those rules? One answer is provided by the COS conformance tests. This battery of tests attempts to insure that a vendor's product correctly implements the protocols. Although COS is not the only creator of OSI tests in the world, it is surely one of the most important. Implementations which pass the appropriate COS tests can receive the COS seal of approval, officially called the COS Mark. Although the COS Mark is currently available for only a few of OSI's many protocols, the intent is to provide a COS Mark for a wide range of OSI products. Like MAP/TOP and GOSIP, the protocol options and interpretations used by the COS conformance tests are currently based on the NIST OSI Implementors Workshop Stable Agreements. When ISPs become available, COS's intent is to base the tests on these international agreements.

SPAG

The *Standards Promotion and Application Group* (SPAG) is a European organization with some strong similarities to COS. Like COS, it is a membership organization. Also like COS, it is sponsoring the development of conformance tests for OSI products.

Recently, COS and SPAG have begun to work together more closely to avoid duplication in the time-consuming work of developing conformance tests.

CEN/CENELEC

In English, this acronym stands for *European Committee for Standardization/European Committee for Electrotechnical Standardization*. As their names suggest, these two bodies produce European standards (called ENs or ENVs) for a variety of areas. Membership is restricted to European nations, primarily those in the *European Economic Community* (EEC). Prior to the creation of EWOS, CEN/CENELEC was a chief sponsor of the development of OSI implementation agreements (which they called functional standards).

Most OSI standards take longer to produce than most of us would wish. This has been especially true in one critical area: *network management*. The best estimates suggest that products based entirely on International Standards for management will not be available until the mid-1990s.

The OSI Network Management Forum

This delay, combined with the perceived market demand for multi-vendor network management products, has led to the creation of the *OSI Network Management Forum*. The Forum was originally established by AT&T, Hewlett Packard, Unisys, and several other companies. Its membership has since increased to include most of the major (and minor) players in telecommunications markets, including IBM and DEC.

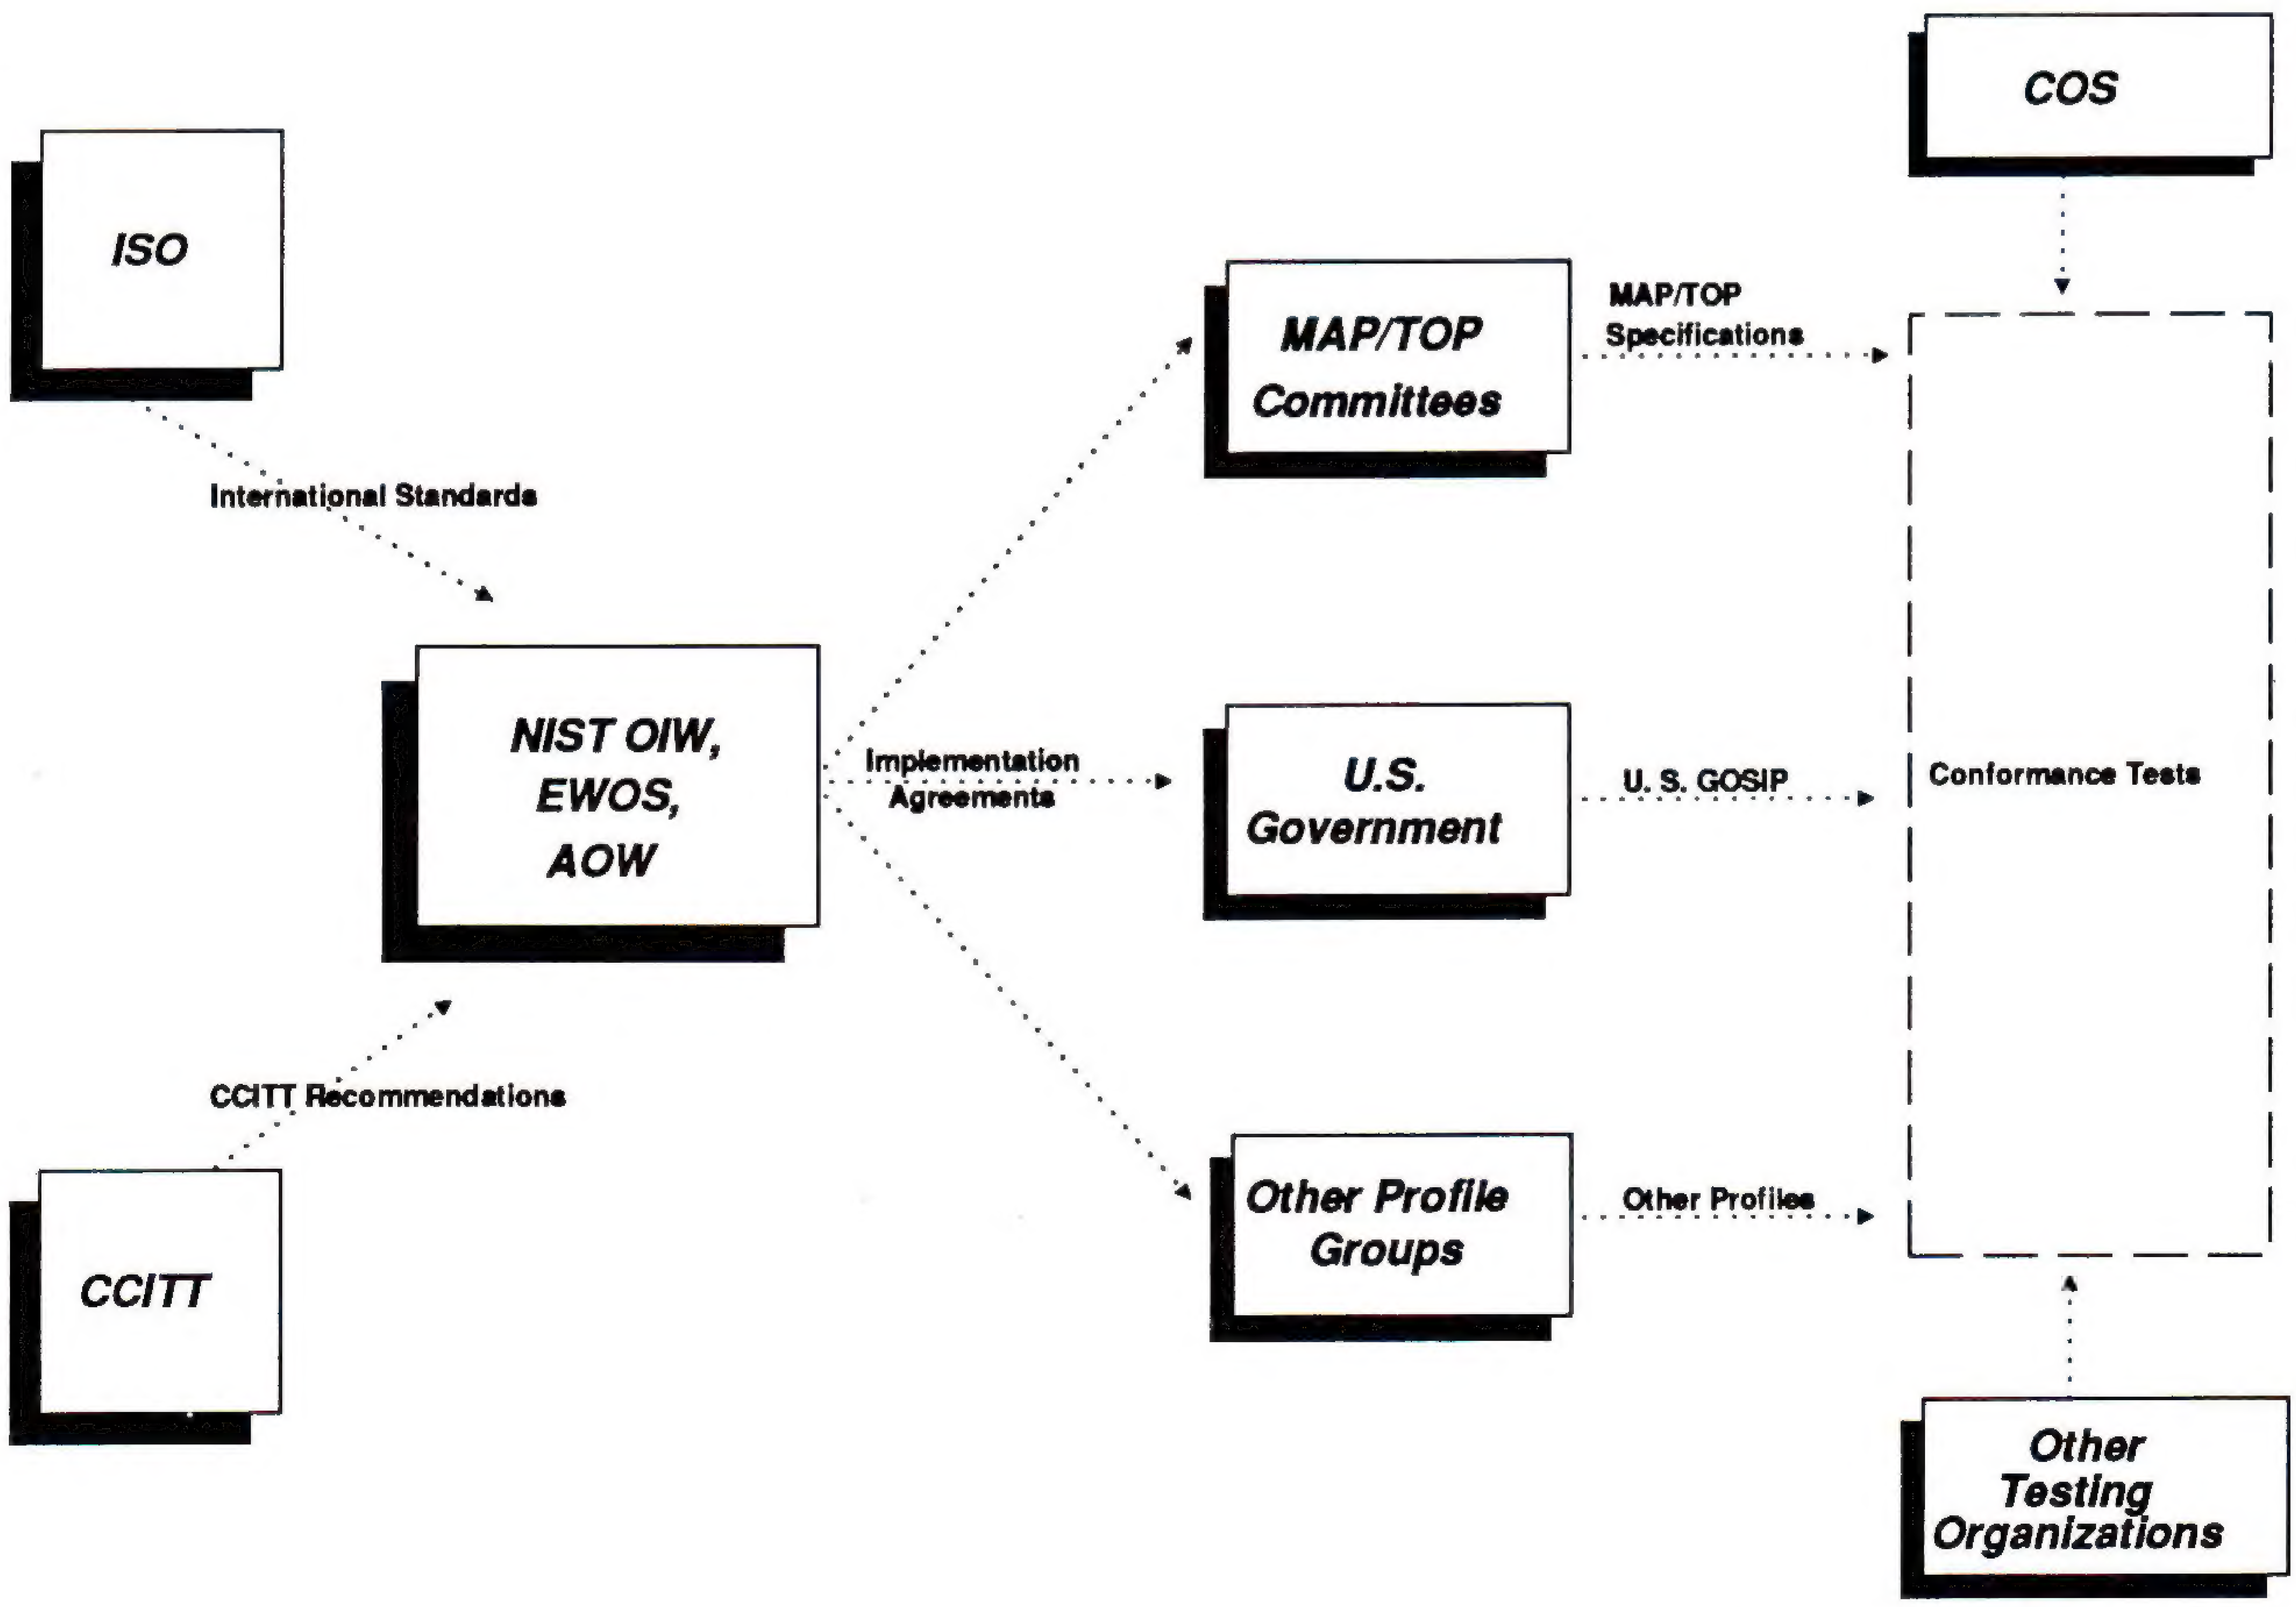

Figure 1: Relationships among some OSI-related organizations

continued on next page

A taxonomy of the players *(continued)*

The goal of the Forum is to create interim agreements for OSI network management. Those agreements will be based on whatever OSI standards exist, but will also include original work developed by the Forum where no standards are yet completed. Since it is not bound by the formalities of the international standards process, the Forum can produce completed agreements much more quickly than can any official ISO body.

Conclusion

The world of OSI is not neat. It is made up of many different overlapping, cooperating, and sometimes competing organizations. The path from idea to standard to tested product is long and complicated (the main elements of this path are shown in Figure 1).

The development of standards for OSI has been something of a departure from the norm for standards bodies. In the past, well understood technologies were standardized, with the standards bodies sometimes just putting their stamp of approval on accepted commercial practices. With OSI, however, the standardizers found themselves out in front, leading rather than following commercial development.

While this approach is probably required if international standards for computer communications are to be achieved, it also has some drawbacks. Among them is this: solving complex technical problems, such as those found in computer communications, is intrinsically hard. Solving those problems effectively for a multi-vendor environment is even harder. Solving those hard problems in the context of the international organizations described above, with all of the accompanying political baggage, as well as reaching agreement among the very wide range of participants, approaches the impossible. Nevertheless, in spite of the technical problems, in spite of the interpersonal and political problems, and in spite of the improbability of perfect results, progress continues to be made towards the noble goal of Open Systems Interconnection.

DAVID CHAPPELL has been active in the design and implementation of OSI protocols for the past several years as a software engineer with NCR Corporation and Cray Research. He has also been an active participant in the NIST OSI Workshops, and currently chairs the Workshop's Upper Layers Special Interest Group. David holds an M.S. in computer science from the University of Wisconsin, and now spends most of his time writing, consulting, and teaching about OSI and related topics.

Articles in our series *Components of OSI* to date:

ISDN	April	1989
X.400 Message Handling System	May	1989
X.500 Directory Services	June	1989
The Transport Layer	July	1989
Routing (ES-IS and IS-IS)	August	1989
The Session Service	September	1989
CLNP	October	1989
The Presentation Layer	November	1989
A taxonomy of the players	December	1989

Still to come in 1990: FTAM, VTP, The Application Layer Structure, Security, ODA/ODIF, X.25, CO–CL Interworking, ASN.1 and more.

A Letter to the Editor

Dear Ole,

This letter is intended to correct the misunderstanding that may have been caused by Vint Cerf's reference, in his 9/89 *ConneXions* letter, to a new OSI Transport protocol class. I encountered some interesting questions at INTEROP 89 from people who wanted to know why on earth ISO, which already has too many Transport protocol classes, intended to define a new one!

Vint's letter refers to the existence of a "fifth class" of OSI Transport protocol. Strictly speaking, of course, there are already five classes (numbered 0 through 4), but it is clear that Vint is talking about a new (sixth) "Class 5," which would target high-performance networks. However, although both ISO and ANSI are interested in the Transport layer implications of high-performance networks, there is no Class 5 Transport, nor does ISO intend to define one.

ISO is working on a number of enhancements to Transport, such as larger Maximum TPDU sizes, finer maximum TPDU size granularity, and selective acknowledgement, which are intended to improve the performance of the protocol. These will become part of some future version of ISO 8073. They will all be compatible, within the existing rules of Transport function negotiation, with the current version of the protocol, and will not involve the creation of a new protocol class.

ANSI-accredited task group X3S3.3 has gone somewhat further, establishing a study project to determine what (if anything) can be done in the Network and Transport layers to complement the emergence of very high-performance subnetwork technologies. The need for such a project was first asserted in X3S3.3 by Dr. Greg Chesson, which perhaps explain the reference to XTP in Vint's letter.

ISO is also preparing a new edition of the Transport protocol specification (ISO 8073), which it expects to publish next year. This new edition will incorporate the text changes that have resulted from the resolution of the defect reports that have been submitted since the standard was published in 1986; the same text changes have already been incorporated into the corresponding 1988 CCITT Recommendation X.224. No new capabilities will be defined in this edition, which does not account for any of the work on enhancements described above.

Cordially,

A. Lyman Chapin
X3S3.3 Chairman

Thanks to Mr. Chapin for the clarification. Keep those letters coming! —*Ed.*

LAN Packet Demultiplexing
or
How You Can Run All Those Different Protocols on One Cable

by James VanBokkelen, FTP Software

Introduction

Most LAN administrators know quite a bit about LANs, such as how data is broken up into packets, and how each host has a hardware address by which it is known to others. Many could outline the basic mechanisms of Ethernet or Token Ring, and lay out the wiring for an office. However, if faced with the question "We've got this Brand X LAN cable all over the place, can I connect my two Brand Y machines to it?," they may well be at a loss for an answer.

Diversity may cause problems

As long as hardware interfaces for whatever LAN standard is in use (Ethernet, 802.5, Starlan, etc.) are available for the new machines, and they run the same version of the same network protocol (Netware, DECNet, TCP/IP or whatever), compatibility is ensured (unless one implementation has bugs). However, this isn't always the case. One reason is that different manufacturers all too often support different LAN protocols: Another department might choose a different LAN O/S for their PCs, or bring IBM compatible machines into a previously all-MAC environment. Another reason is that LAN protocols are frequently specialized: the original installation might be using a PC LAN O/S to share files and printers, and the new users might need to use TCP/IP over the LAN to log in to a mini-supercomputer and run linear programs or CAD software. They can't use the LAN O/S because its developers didn't include remote login in their design.

Multiplexing

Fortunately, the designers and implementors of most of the major LAN standards had *multiplexing* in mind. LAN cabling is expensive, and LAN bandwidth is usually much larger than any single host could use, so it is a natural reaction to make the resource as widely shareable as possible. The first level of sharing allows many different clients and servers using protocol A to use the same cable, at the same time. The second, which is designed into all but a few LAN standards, allows the machines speaking protocol A to share the cable with others using protocols B and C, and even allows sophisticated machines to use several different protocol families at once.

How does this multiplexing work? When you look at how data moves over the LAN media, you find that all LAN technologies break each machine-to-machine data flow up into a series of separate messages, or packets. Each packet has a source and destination address, and framing to separate it from other packets. This allows many different hosts to share a common wire to communicate, without the overhead and cost of setting up separate circuits between each pair of hosts.

It also makes it easier to design higher-level protocols that guarantee delivery of data even in the face of transmission errors. Thus, the first level of LAN multiplexing is based on packetizing and individual machine addresses. The recognition of hardware addresses and framing/unframing packets is almost always done in hardware, to reduce the load on the host CPU.

To implement the second level of multiplexing requires some way of telling protocol A packets from protocol B packets. At first, this might not seem necessary, unless some of the machines want to speak several different protocols at once, because hardware addresses would seem to ensure that packets from protocol A are only received by machines that speak A. However, almost all LAN protocols make some use of broadcast packets, sent to a special hardware address which all machines receive. This is most often done to locate machines whose hardware address isn't known, or to locate a server of some sort. Since hardware addresses are usually assigned arbitrarily, as interfaces are manufactured, and change when interfaces are replaced, broadcast protocols are very attractive compared to long tables of hexadecimal numbers on each LAN node.

Ethernet Packet Type

There are several different second-level demultiplexing mechanisms: one of the earliest, and still the most widely used, was that selected by DEC, Intel and Xerox when they defined *Ethernet*. Each DIX (or Bluebook) Ethernet packet has a 16-bit "Ethernet packet type" field in the header immediately after the destination and source addresses. All XNS packets have an Ethertype value of 0x600, IP is assigned Ethertype 0x800, DECNet Maintenance & Operations Protocol uses 0x6002, and so forth. Every DIX Ethernet protocol, from widely used public specifications like IP to private protocols used by a single product from a single manufacturer, must have a unique number assigned to it, so it is a good thing that there are 65,000 available.

Ethertypes were initially allocated by Xerox, but the IEEE has taken over that function. (Regrettably, neither organization has ever made the list of assigned types public, which sometimes complicates Ethernet debugging. Several different lists of known Ethertypes exist, maintained by various commercial and educational network sites, but they are necessarily incomplete.)

IEEE Packet Header

The second most widely-used demultiplexing mechanism is the IEEE packet header format specified in 802.2. 802.2 headers are in use on 802.3 (IEEE Ethernet), 802.4 (Token Bus), 802.5 (IBM-type Token Ring) and FDDI networks. In each case the 802.2 header immediately follows the physical-layer header, which varies according to the media in use. The 802.2 header uses the same amount of space (16 bits) for demultiplexing information as DIX Ethernet, but the IEEE chose to divide it up differently. The first byte is the *Destination Service Access Point* (DSAP), and the second is the *Source SAP* (SSAP), where SAPs roughly correspond to Ethertypes. Because a number of bits have been reserved to define special SAPs, only 24 unique values are available for demultiplexing different standard protocol formats.

The 802.2 header is normally either 3 or 4 bytes long, depending on which IEEE Logical Link Layer Control method is in use. LLC1 uses 3 bytes, and provides a simple datagram service. LLC2 uses 4, and provides reliable link-level delivery of messages. The LLC2 scheme is attractive to designers of small-scale LAN protocols, because it can greatly simplify a lightweight transport layer. However, LLC2 is much less useful in internetworking, where many different LAN and WAN segments are connected by packet routing nodes.

continued on next page

LAN Packet Demultiplexing *(continued)*

Most implementers of large-scale protocols have chosen to use LLC1 for its lower overhead, since the presence of different media types and routers requires a separate end-to-end reliable delivery mechanism even if LLC2 is supported on all links for all paths.

OUI

Because the number of available SAPs is so small, the IEEE has been very reluctant to assign values to protocols it feels are unimportant, or restricted in scope. While several vendors have simply appropriated SAP values for their private protocols, others have made use of the Sub-Network Access Protocol within 802.2. This is identified by a reserved SSAP/DSAP value, and adds a 3 byte *Organizationally Unique Identifier* (OUI) to the normal header.

The TCP/IP community obtained an OUI of 0x000000, and defined it as being followed by a 16-bit DIX Ethertype, in order to provide access to any existing Ethernet protocol on all 802-type LAN media. The eight bytes of demultiplexing information this represents is the largest in common use today.

Other LAN media may use 802.2 headers, or type fields ranging from 8 to 32 bits, but the principle is the same; soon after a packet is received, low-level software can look at some well-defined part of the packet header to find out which protocol it is destined for. If this host has no higher-level software to deal with a particular packet, it should be discarded quickly, so as not to waste CPU cycles.

Vendors

Of course, as everyone who has ever used a computer knows, vendors are rarely perfect. In the LAN world, there are certainly careless vendors, ignorant vendors and self-centered vendors. The careless kind can be recognized when their software hangs or misbehaves when confronted with a packet it doesn't understand, or when a machine crashes and jams the LAN cable until it is powered off. The ignorant didn't fully understand the standard they implemented, and send packets that violate the spec, or ship hardware that can only talk with each other and not with other vendors' cards.

Self-centered vendors are sometimes the hardest to cope with. Using an un-allocated packet type or SAP value can lead to protocols that won't share a cable or an interface. Implementing part of a standard but discarding the rest as inefficient has been known to cause conforming implementations to report an error for every packet they receive. Perhaps the most common instance of self-centeredness is abuse of broadcast facilities, producing systems that don't scale up well. Frequent broadcasts are OK in a cozy little community of 20 personal computers, but not when you have 400 nodes representing 5 departments on the cable, particularly when someone suggests a per-packet charge to "allocate overhead costs better."

This presents a dilemma to the LAN designer or administrator: If all the different protocols abide by the rules, you can get a lot more return out of an expensive cabling job, improve interdepartmental communication and use specific networking protocols for specific jobs. If they don't, you get headaches, and unhappy users.

My opinion is that even if you can segregate each function right now, it won't last. Somebody will need to do their next CAD project on a graphics workstation instead of a PC, but they'll want to send their finished drawings to the PC users' plotter. Purchasing, Manufacturing, Engineering and A/P will want to share an inventory database. Somebody will notice the cost of the tape drive to back up each individual minicomputer.

Given this, any resource that helps you design, purchase, install and maintain large, heterogeneous networks is very valuable. Books, courses, conferences and seminars each have their place, but the most powerful resource is other LAN-using sites.

Staying informed

If you can get access to one of the national news/mail/conferencing networks, either commercial, like *CompuServe* or *The Source*, or non-commercial like USENET, do so. Find the LAN-related forums with high signal-to-noise ratios and post your questions, answering in turn when others want information you have. In addition, or if you can't do it electronically, check out local computer groups, go to vendor's user-group meetings, ask salesmen for references and develop your own network of resources. The diversity of the LAN-using community practically ensures that someone else has experience with the hardware or software combination that you're interested in.

JAMES VanBOKKELEN's undergraduate relationship with MIT ended (scoreless tie) in 1980. From 1980 to 1985, he was Manager, Software Development for Perception Technology Corp., working on touch-tone data entry/voice response devices (does the IRS "Teletax" system ring a bell? No? Oh, well.) In 1986, he helped found FTP Software Inc., becoming the company's first VP of Marketing. Circumstances being what they are, he found himself getting sucked back into Software, replacing John Romkey as VP in 1987, and on the principle that the biggest pieces rise to the top, replaced Roxanne VanBokkelen as President in 1988. Most of his time is spent working on projects no one else here will touch, like the IETF Host Requirements Working Group, RFC 1001/1002 NetBIOS, TCP subtleties, and answering lots of e-mail.

A Clarification

The November 1989 issue of *ConneXions* has a map of the INTEROP Show and Tel-Net showing links to the outside world. The T-1 circuit from the shownet terminates in something labelled "FEBA West." Several of our readers have asked what the acronym FEBA stands for, so we decided to pass the question on to Milo Medin of NASA Ames Research Center. Milo replies:

Well, when those of us who run the backbone networks were discussing the configuration at NASA Ames, the name "Inter-agency Interconnect Network" was a bit of a mouthful. Since this piece of Ethernet has some of the most complicated routing in the Internet being performed on it, I likened it to the front line of routing, and suggested we call it a FEBA (pronounced *fee bah*), which is a military term for the *Forward Edge Battle Area* (aka the front line). People liked it so much, it just sort of stuck. It's actually the official name now.

Thanks,
Milo

Information Sources

We at ACE are often asked about sources for information on TCP/IP OSI and other networking topics. In this article we will give an overview of what is available. The article was compiled in part with the help of the USER-DOC Working Group of the IETF. (See page 1)

Books

There are many good books on telecommunications in general and computer networking in particular. The following is a list of some of the relevant books. Several of them have been reviewed in previous issues of *ConneXions*. The year/month the review appeared is listen in parenthesis.

- *Computer Networks, Second Edition* (1/89)
  by Andrew S. Tanenbaum, Prentice-Hall, ISBN 0-13-162959-X.

- *Internetworking With TCP/IP—* (1/89)
  Principles, Protocols, and Architecture
  by Douglas E. Comer, Prentice-Hall, ISBN 0-13-470154-2

- *An Introduction to TCP/IP* (4/89)
  by John Davidson, Springer-Verlag, ISBN 0-387-96651-X

- *Handbook of Computer-Communication Standards*
  Volume 3: Department of Defense Protocol Standards
  by William Stallings et al, Macmillan, ISBN 0-02-948072-8

- *Internetworking Computer Systems* (10/89)
  by John McConnell, Prentice-Hall, ISBN 0-13-473976-0

- *Internetworking—A Tutorial on Bridges and Routers*,
  Available from 3Com Corp., Santa Clara, CA, 408-562-6400.

- *Internetworking—An Introduction*, Available from
  The Wollongong Group, Palo Alto, CA, 415-962-7100.

- *Innovations in Internetworking*
  edited by Craig Partridge, Artech House, ISBN 0-89006-337-0.

- *!%@:: A Directory of Electronic Mail Addressing and Networks*
  by Donnalyn Frey and Rick Adams, O'Reilly and Associates, ISBN 0-937175-93-0.

- *The Matrix: Computer Networks and Conferencing* (11/89)
  Systems Worldwide by John S. Quarterman, Digital Press
  ISBN 1-55558-033-5.

- *Users' Directory of Computer Networks*
  by Tracy Lynn LaQuey Office of Telecommunication Services, University of Texas System. Soon to be published in book form by Digital Press.

- *The Open Book: A Practical Perspective on OSI* (11/89)
  by Marshall T. Rose, Prentice-Hall, ISBN 0-13-643016-3.

RFCs and Military Standards

Request For Comments (RFCs) contain all the protocol specifications used by the Internet community. Note that only some of these are "Official Internet Protocols". (The list of official protocols is itself an RFC and is published periodically). See also RFCs 999, 1000, and 1012 for a guide to RFCs.

There are five (5) Military Standard (MIL-STD) documents. The IP and TCP MIL-STDs are completely different from the corresponding RFCs yet intending to define the same protocols. The FTP, MAIL, and Telnet MIL-STDs are copies of the corresponding RFCs. Implementors should always use all the available information. There are some known errors in MIL-STDs (specifically, see RFC 963 and RFC 964). In cases of confusion, the RFCs should take precedence over MIL-STDs.

Getting RFCs

RFCs are available from the DDN Network Information Center at SRI International. Hardcopies can be ordered by calling 800-235-3155 or 415-859-3695. Online versions may be copied via anonymous FTP from the RFC: directory on host NIC.DDN.MIL (10.0.0.51 or 26.0.0.73). (Check out the file RFC:RFC-INDEX first). If you are outside the "connected Internet" and don't have access to FTP, you can still receive RFCs by simply sending electronic mail to SERVICE@NIC.DDN.MIL. Put the RFC number in the "Subject:" field (e.g. Subject: RFC 980). The DDN NIC also publishes the *DDN Protocol Handbook*, a four-volume set, which contains most of the important RFCs related to TCP/IP. Although parts of this document are now a little out of date it is definitely worth getting for reference.

Getting MIL-STDs

MIL-STDs are available from:

- Naval Publications and Forms Center, Code 3015
  5801 Tabor Avenue
  Philadelphia, PA 19120
  215-697-3321 (order tape) 215-697-4834 (conversation)

Basic Beige RFCs

Joyce Reynolds of USC-ISI recently compiled this list of "Basic Beige" RFCs. Note that this list will change from time to time as new RFCs are added. In particular the Official Protocols and Internet Numbers are issued periodically, make sure you obtain the latest version.

RFC 768	User Datagram Protocol (UDP)
RFC 791	Internet Protocol (IP)
RFC 792	Internet Control Message Protocol (ICMP)
RFC 793	Transmission Control Protocol (TCP)
RFC 821	Simple Mail Transfer Protocol (SMTP)
RFC 822	Standard for the Format of ARPA Internet Text Messages
RFC 826	Ethernet Address Resolution Protocol
RFC 854	Telnet Protocol
RFC 862	Echo Protocol
RFC 894	A Standard for the Transmission of IP Datagrams over Ethernet Networks
RFC 904	Exterior Gateway Protocol
RFC 919	Broadcasting Internet Datagrams
RFC 922	Broadcasting Internet Datagrams in the Presence of Subnets
RFC 950	Internet Standard Subnetting Procedure
RFC 951	Bootstrap Protocol (BOOTP)
RFC 959	File Transfer Protocol (FTP)
RFC 966	Host Groups: A Multicast Extension to the Internet Protocol
RFC 974	Mail Routing and the Domain System

continued on next page

Information Sources *(continued)*

RFC 1000	The Request for Comments Reference Guide
RFC 1009	Requirements for Internet Gateways
RFC 1010	Assigned Numbers
RFC 1011	Official Internet Protocols
RFC 1012	Bibliography of Request for Comments 1 through 999
RFC 1034	Domain Names—Concepts and Facilities
RFC 1035	Domain Names—Implementation
RFC 1042	A Standard for the Transmission of IP Datagrams over IEEE 802 Networks
RFC 1048	BOOTP Vendor Information Extensions
RFC 1058	Routing Information Protocol
RFC 1065	Structure and Identification of Management Information for TCP/IP-based internets
RFC 1066	Management Information Base for Network Management of TCP/IP-based internets
RFC 1084	BOOTP Vendor Information Extensions
RFC 1087	Ethics and the Internet
RFC 1095	The Common Management Information Services and Protocol over TCP/IP (CMOT)
RFC 1098	A Simple Network Management Protocol (SNMP)
RFC 1101	DNS Encoding of Network Names and Other Types
RFC 1112	Host Extensions for IP Multicasting
RFC 1117	Internet Numbers
RFC 1119	Network Time Protocol (NTP)
RFC 1122	Requirements for Internet Hosts—Communication Layers
RFC 1123	Requirements for Internet Hosts—Application and Support
RFC 1130	IAB Official Protocol Standards

List of implementations

The DDN Protocol Implementations and Vendors Guide, also available from the DDN NIC, contains listings of vendor products for TCP/IP, X.25, and (a few) OSI products. This document is extremely useful for finding protocol software for your particular system.

OSI documents

CCITT and ISO documents may be ordered from:

- Omnicom Inc.
  115 Park Street, SE
  Vienna, VA 22180-4607
  703-281-1135

Network Information Centers

Most of the the larger networks operate *Network Information Centers* (NICs) which provide help and information for current or potential users. The major NICs are listed below.

- DDN Network Information Center
  SRI International
  333 Ravenswood Avenue
  Menlo Park, CA 94025
  800-235-3155 or 415-859-3695 NIC@NIC.DDN.MIL

- NSF Network Service Center
  BBN Laboratories
  10 Moulton Street
  Cambridge, MA 02238
  617-873-3400 NNSC@NNSC.NSF.NET

- CSNET Coordination and Information Center (CIC)
  BBN Laboratories
  10 Moulton Street
  Cambridge, MA 02238
  617-873-2777 CIC@SH.CS.NET

- The BITNET Information Center
  EDUCOM, Suite 600
  1112 16th Street, N.W.
  Washington, DC 20036
  202-872-4200 BITNET: INFO@BITNIC

Mailing lists

Electronic mailing lists exists for both TCP/IP and OSI. The ones of most interest to readers of *ConneXions* are probably "ISO@NIC.DDN.MIL" and "TCP-IP@NIC.DDN.MIL". These lists are also distributed on USENET and are available as **comp.protocols.tcp-ip** and **comp.protocols.iso** respectively.

Bibliography

The Experimental Literature of The Internet: An Annotated Bibliography, Jeffrey C. Mogul, Digital Equipment Corporation, WRL Technical Note TN-1, August 1988. Available from:

Technical Report Distribution
DEC Western Research Laboratory, UCO-4
100 Hamilton Avenue
Palo Alto, CA 94301
E-mail: WRL-Techreports@decwrl.dec.com

Online documents

A number of informative articles are available over the Internet. Some of these online files are listed below.

Introduction to the Internet Protocols by Charles L. Hedrick, Rutgers University. Available from host `topaz.rutgers.edu`, in the directory `pub/tcp-ip-docs`, filenames `tcp-ip-intro.1` and `tcp-ip-intro.2`.

The Hitchhiker's Guide to the Internet by Ed Krol, University of Illinois. Available on host `NIC.DDN.MIL`, in the `RFC:` directory, filename `RFC 1118.TXT`.

Network Reading List by Charles Spurgeon, The University of Texas at Austin. Available on host `emx.utexas.edu`, in the directory `pub/netinfo/docs`, filename `network-reading-list.txt`.

Glossary of Networking Terms by Charles Spurgeon, The University of Texas at Austin Network Information Center. Available on host `emx.utexas.edu`, directory `pub/netinfo/utnet`, filename `gloss.txt` or `gloss.ps`.

Tutorials and seminars

Courses on TCP/IP and OSI are available from Advanced Computing Environments. Tutorials are offered several times a year, in conjunction with our *INTEROP*™ conferences, as stand-alone Tutorial Series, or as in-house tutorials tailored to your specific needs. Additionally, we schedule seminars on specific topics such as TCP Performance and Internetwork Routing.

Periodicals

Newsletters such as *ConneXions* provide a good way to keep informed on a variety of networking topics. Most of the national and regional networks have their own newsletters. There are many academic journals dedicated to networking, for instance ACM SIGCOMM's *Computer Communications Review*, IEEE's *Network*, and North-Holland's *Computer Networks and ISDN Systems*.

CONNEXIONS
480 San Antonio Road
Suite 100
Mountain View, CA 94040
415-941-3399
FAX: 415-949-1779

FIRST CLASS MAIL
U.S. POSTAGE
PAID
SAN JOSE, CA
PERMIT NO. 1